

**THE MINOS TIMING SYSTEM:
TIMING DISTRIBUTION (FAR DETECTOR ONLY)
AND GPS TIMING**

C.Perry - Oxford – 01jun00

**CHANGES SINCE DEC99 REVIEW
AND RESPONSES TO COMMENTS**

1. Scope of Document

This note outlines the changes from the proposal as previously reviewed, and documented in:

The MINOS Timing System: Timing Distribution (Far Detector Only) and GPS Timing:

Part 1: Timing System Principles. C.Perry, 07dec99

Part 2: Timing System Hardware. C.Perry, 07dec99

The MINOS Timing System: A Note Concerning the Fundamental Options.

C.Perry 10dec99

The system remains generally as described. Decisions have been made on the options discussed, and there have been some minor technical changes.

2. System Decisions

a) Scope of System

The system continues to cover timing distribution at the far detector only. A desire that the near detector timing be locked to the Main Injector RF (which is swept in frequency over a significant range, and behaves in a complex fashion), and lack of time in the schedule to explore all the ramifications of this, made it problematic to require the system to cover both.

b) Receiver Location

The GPS receiver will be underground at both locations, fed by a fibre optic antenna link. This was the simplest to install and understand, and required the least development. It also avoids a problem that had not been appreciated, that there are to be no cables installed from the Soudan surface building to the mineshaft. The antenna and head end of the link can go on the existing mine buildings, which was unlikely to have been possible for a bulkier installation otherwise needed. There appear to be no objections to the fibre optic link.

c) Far and Near Operating Modes

The system at the far detector will distribute the GPS time provided by the receiver. This is simplest, both in concept and implementation. The near detector will be different; given that it may be required to use a clock swept with the MI RF, this is almost unavoidable. Instead, at the near detector the GPS system will timestamp signals from the detector timing system. This event timing function is an option in the GPS receiver which it is intended to use. In the unlikely case that the cost is excessive, a timing system Central Unit as used at the far detector could be programmed to provide event timing.

d) High Performance Clock

A high performance local clock will not be included. It is considered that in the unlikely event of the loss of the GPS input, the fairly high grade crystal oscillator in the GPS receiver will maintain adequate timing until the input is restored. However, the receiver which it is intended to use can optionally be upgraded with a rubidium clock if this ever proves necessary.

e) Buffer Changeover Signal & VME Interrupts

A signal to change between buffers in the front end modules is now to be generated centrally and distributed over the timing system. Changeover will occur at the second marks and at intermediate times. This has no impact on the hardware, and little on the coding. There will now be two types of time marker code transmitted, one for seconds and one for the intermediate buffer changes. It will be a fourth LVDS signal distributed across the front end crate J3 position backplane.

And at the emphatic request of the people at RAL, a VME interrupter capability has been added to the receiver module. This is so the signal can also be passed to the processor board. There were apparently difficulties in a direct electrical connection to an input on the board. A full VME interface will still not be provided.

Implementation Changes

a) Network Time Server

The TrueTime XL-DC GPS receiver, which is our intended unit, can now have a Network Time Server option fitted. This gives the simplest solution, with entirely standard commercial hardware. It will be adopted unless the cost proves prohibitive.

b) Scheme of Development

The original scheme called for handwired prototypes, before going to pcbs, and this influenced the design. Schedule and requirement changes made this look unattractive, and we decided instead to go straight to a 'test bed' pcb.

This has now been done, and includes all the functionality that is required for both central unit and receiver module. It should be possible to completely develop the logic with this board, and to use it for the 'vertical slice' tests. It is a 6U board, that will fit in a standard 6U VME crate (allowing the VME interrupter to be checked), or in low position in a 9U (allowing it to drive the front end cards). The size is small enough to go in the Central Unit, and to go through our reflow oven (so we don't have to subcontract assembly), and to be reasonably cheap.

We will then do a revised version if necessary, to go in the Central Unit. We would probably not strip off unwanted capabilities, merely part assemble the boards.

And we will do a 9U version as the Receiver Module. This is just a matter of enlarging the board outline, and moving the VME interrupter up and the optical input to the front. The changes have been anticipated: critical signal lines already include terminating resistors. A good deal will probably be stripped off in this case. Consideration will be given to reducing it from its present four layers to two, to save cost, but this seems problematic.

c) Central Unit Functionality

The Central Unit no longer has to provide the fibre optic links to a GPS receiver on the surface, so some of the ancillary io functions are not required. However, similar and even more extensive capabilities are being included on the board to allow for unanticipated requirements. The board having been fixed for other reasons as 6U VME in size, there is little cost in this. This also extends to the VCSEL outputs: allowance is made for four.

d) Logic

We have upgraded the logic from a PLD to fairly large FPGA. This was done to secure more flexibility. It would, for example, allow event timing to be added in the central unit. Both central unit and receiver module will have the same type of FPGA (XCV50, though a larger may be fitted for development work), being derived from the common development board. The cost penalty of excess capability on the receiver seemed insignificant set against the savings in development.

e) Clock Multiplication

This is needed because the timing receiver output of 10MHz has to be converted to the 40MHz required. It was originally done with a VCO and external logic, and included only on the receiver.

This has been changed to use a Cypress clock chip (second sourced by Pericom), with a self-contained PLL. This had been recognised as a simpler solution but was considered to need further investigation: it is now felt to be suitable. The original approach could still be fitted on the development board if there are problems.

f) Optical Carrier Frequency

A carrier frequency of 2.5MHz, with pulse width modulation superimposed, had been proposed. This will now be increased to 10MHz.

This can be done with no penalty because of the way the clock multiplication is now being done, and the development scheme adopted. It means there is no penalty in multiplying the input clock up to 40MHz in the central unit, so the higher carrier frequency can as easily be used.

At the same time, the wider loop bandwidth of the clock multiplier chip being used in place of a VCO in a custom loop necessitates minimal interruption of the reference input while an execute signal is being received; the higher clock is then advantageous.

g) Optical Coding Scheme

This will be changed somewhat. It was originally somewhat constrained by the limited logic possible in the PLDs. The new scheme will be similar, but with better error handling, some added capabilities, and an extensible design. Details remain to be defined.

h) Fine Timing Generation for Execute Signal

The original approach used an Analog Devices delay generator chip to provide fine subdivision of the 10MHz clock period. This was so a finely adjustable execute signal could be generated to allow the front end fine timing system to be thoroughly checked. This method has been dropped, as the new design allows easier ways that do not require any setting up

The basic one is to use the capabilities of the Cypress clock chip. This employs a multistage ring oscillator, with adjustable stage delay, as its VCO. The chip provides four outputs, that can be derived from different taps on the oscillator. This allows a second clock output to be produced adjustable in phase in steps under 1ns. We would derive the execute from this second clock.

An alternative, that we will use if the Cypress chip proves to have problems, uses the DLLs in the Xilinx FPGA. These allow us to multiply the clock and get timing edges 3ns apart. This is not fine enough, but analog interpolation to 1ns (and possibly finer) is trivial: an external RC network driven by two outputs and feeding back to a third is all that is needed.

i) Clock Distribution in the Front End Crate

The Harvard group has defined this as using four LVDS signals over a J1 type VME backplane used in the J3 position. They are using pairs of signal lines, driven by the receiver board at one side of the crate, and terminated at the other.

This seems perfectly adequate, with the reservation that there may be significant irregularities due to connectors and stub traces on the receiver boards given the speed at which LVDS (and the FPGAs being fed by it) can operate.

Any such problems are easily avoided by slowing the edges put on the backplane by the LVDS drivers. All this takes is a small capacitor across each output. I have provided this on the development board, with 18pF fitted for the backplane tests. This gives about a 5ns risetime on the backplane, which is more than adequate and will not contribute more than 0.5ns skew between different crates.

Responses to Comments (list from Alfons)

a) Earl Peterson, Adam Para: too many alternatives

b) Adam Para: too little prototyping

Agreed. The reason was that MINOS requirements kept changing radically.

It was only last August the near detector was changed to meet the requirements of single turn extraction by using the QIE chip. This appeared to mean using a 53MHz clock. Since I would have to make a clock system for this, there was no point in continuing with the CERN TTC solution that I had been working on, since that was tied to 40MHz. I developed an electrical scheme to handle both frequencies, which I documented for the revised TDR electronics chapter, and was planning to prototype for the last review.

At that point I was told to do a system for the far detector only. The electrical system was no longer optimal, and I designed a new optical system. It would have been possible to revert to the CERN TTC system, but this was rather complex, there were problems with their schedules, and with the system architecture changed so there were only 16 destination crates the advantage of having a receiver ASIC were minimized. I had intended to prototype the new system for the review, but ran out of time.

With the changes to the timing system requirements, and the work at Oxford on the VA chip and the front end board (VFB) that followed the August decisions, work on the GPS part of the timing system had to be postponed as less critical, being commercial equipment and not useful until the near detector also was operational. This left quite a lot of options undecided. And since the questions had never been discussed by the Collaboration, it seemed useful to have them aired at the review.

The outcome of the review was that it appeared the system would once again be required to do both detectors. I stopped work on it, waiting for a decision. Finally it was decided at the February meeting that it would in fact still be for the far detector only.

The schedule now was very tight, given that clock distribution hardware would be wanted for the vertical slice test in October or thereabouts. A final redesign was done that eliminated the separate prototyping phase, minimized the overall amount of work that had to be completed before the vertical slice test, and added as much flexibility as I could to allow for requirements that kept changing.

This has meant that I now have a board which is close to the final design of both central unit and receiver module, and should be able to function in both positions for the vertical slice test. It is however not yet fully assembled and functional, so the system is still not prototyped.

The pressure to complete this board has meant that I have still only been able to do a limited amount on the GPS part. We have however learnt enough to be able to reduce it to a single proposed scheme. The only issue with this is if the cost is sufficiently over my previous guesses to make us consider an alternative to the TrueTime unit: I don't think this will be necessary.

I quite agree the situation has been entirely unsatisfactory.

c) Pete Border: use optical cable in shaft

Definitely. This was always intended for the mine. And since it is now clear we are using the optical antenna link there, we can and will do the same at Fermilab. The only disadvantage is cost.

d) Robert Hatcher: make near and far GPS the same

Yes. The only exception is that we might not fit the event timing plug-in at the mine, where it isn't needed. The optical antenna link will be the same.

e) Peter Wilson: make near and far as similar as possible

I agree this is desirable, but the degree of similarity has been limited by the decision in February that the near detector clock should be locked to the main injector RF. This is swept in frequency and has phase discontinuities, and there appeared to be other complexities which it would take time to understand.

It was necessary at that point to freeze the system design for the far detector, to be able to meet the schedule. It was decided to make the far detector system as simple as possible, since it was impossible to be sure of correctly anticipating the requirements of the near.

This brings about the situation that the far detector distributes clock signals locked to GPS time, whereas at the near the GPS receiver is likely to be used to timestamp events from the local detector clock system, and thereby allow the local timing to be related to GPS time.

f) Peter Wilson: define the requirements for the timing system

This has never been done. There has never been a definition in any detail from the physicists of what the system is required to do. Equally significantly, there has never been any analysis of how well and how frequently the calibration system will be able to calibrate the clock distribution, which is fundamental to setting its performance requirements.

What I proposed was based on my general understanding of what was required, and my belief that a system of the sort proposed would be both comfortably adequate, and reasonably close to the best possible without special approaches and letting timing skew requirements dominate the entire system design.

I also took into account the design of the rest of the front end. I believe the proposed distribution system has considerably better skew performance than the rest of the front end electronics which it drives.

Here is a rule-of-thumb justification for the design. The optical fanout and fibres contribute delays that are stable to better than 100ps and can be ignored. The delay through the receiver module is about 25ns. The long term variation of this will be under 10%: ie under 2.5ns. The short term variation over a day or so can be expected to be at least an order of magnitude less: ie under 0.25ns. Thermal effects are potentially rapid, but short term temperature excursions at either site should (unless the equipment is being worked on, which can be disregarded as not a normal condition) be under 2°C. Delays

can be assumed to vary by around 0.3% /°C or less, which would give 0.15ns. So we can expect short term stability of the order of 0.3ns or better.

g) Peter Wilson: get GPS receiver underground

Agreed. This we will do, since there seems no problems with the optical link.

h) Peter Wilson: study the performance of the optical link

This is an off-the-shelf option from TrueTime. No-one reports any problems with them; our local TrueTime rep says he has supplied three with no complications. There does not seem anything more we can or need to do.

There are in fact two different links available from TrueTime: their standard commercial one (L1 frequency only), and a military one (L1 and L2) which is somewhat more expensive. The rep is trying to find out if there would be any advantage to us in the second, which is slightly more expensive.

i) Peter Wilson: make an effort to maintain the GPS signal (heated dome)

Agreed. I intend that we give it some attention, see how it goes for the first two winters before we really need GPS, and then upgrade the arrangements if necessary. But the Minnesota climate is not so severe that extreme measures will be needed. I will look for a commercial solution, and if not put the antenna in a small dome with a self-regulating heater.

j) Earl Peterson: define GPS software responsibilities

Refer to Alfons, but I believe this has been addressed. And given that we use the commercial Network Time Server option, the requirements are minimal.

k) Adam Para: don't need GPS

Agreed, at least if we don't start with the low energy beam. But this is a minority opinion.

l) Adam Para: preference for local rubidium clock

I half agree, but with a decent receiver underground the actual impact of occasional loss of lock to GPS doesn't have any significant impact. The TrueTime XL-DC will accept a rubidium clock, but the cost (\$8K) is too high for me to justify. It can however be fitted later as a factory upgrade, provided the receiver is specified for this (it needs a slightly larger chassis).

3. Outline of Schedule and Budget

a) Development Board:

produce pcb	done
test driving backplane	in progress
full assembly of 3 boards	4 wks to week 30 (starts July)
electrical functional test	4 wks to week 34
VCSEL link test	3 wks to week 37
programming operational logic	3 wks to week 40
system test	2 wks to week 42 (ends mid Sept)

- This has functional clock distribution (albeit not in its final packaging) available for the vertical slice test.
- It assumes I (CP) am working on it 80%, with some technician support and no other assistance.
- This period starts after Ely and my work with IDE on the new VA chip, and ends before the new chips are returned for testing.
- Phil Sullivan may be available to do some work, depending on the work needed on the VFB.
- After this the Central Unit, the Receiver Module, and the GPS system proceed in parallel, with my total workload well under 50%.

b) Central Unit

fix pcb	1 wk	to week 43
procure pcbs	2 wks	to week 45
assemble pcb	2 wks	to week 47
obtain hardware	2 wks	to week 49
assemble first unit	1 wk	to week 50
(Christmas/new year break	2 wks	to week 52)
test sample unit	2 wks	to week 02
assemble remaining boards/units	2 wks	to week 04
test units	2 wks	to week 06

c) Receiver Module

derive pcb	2 wks	to week 44
first batch of pcbs (eg 6)	4 wks	to week 48
assemble batch	4 wks	
(Christmas/new year break	2 wks	to week 52)
...		to week 02
test batch	2 wks	to week 04
main batch pcbs		
contract assembly		
acceptance test		

- Production standard central units and receiver modules are shown as available before the end of February 00.
- The remainder of the receiver modules can follow as convenient: they are not required for some considerable time.
- The schedule assumes modest availability of a technician and Phil, and could be accelerated with more resources.
- My involvement is necessarily limited: I am committed to another project.

d) GPS

The GPS hardware procurement will take place during the rest of this year. It is not a resource intensive task, and the only reason to do it on this timescale (apart from convenience) is so we can spend the money before April 01 if necessary (ie in this UK financial year).

Defining the optical fibre requirements is the only thing urgently needed; this will be done in the next month.

e) Budget

A very approximate budget is:

Central units	3 @ £2K	£6K
Receiver modules	25 @ £500	£13K
Distribution fibre cables	16 @ £60	£1K
GPS receivers	4 @ £5.5K	£22K
GPS antennas/links	4 @ £2.3K	£10K
		=====
	Total:	£52K

This is slightly higher than the previous estimate, and a little more than the £45K remaining in the UK budget for the timing system up to April 01. This was intended to cover the entire system. This should not pose any problems; half the GPS expenditure could easily be delayed until next financial year.

CAUTION: these figures are still very approximate. The GPS prices are based on preliminary verbal figures, for roughly the configuration we are interested in.